

This Page Is Inserted by IFW Operations  
and is not a part of the Official Record

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning documents *will not* correct images,  
please do not report the images to the  
Image Problem Mailbox.**

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
29 March 2001 (29.03.2001)

PCT

(10) International Publication Number  
**WO 01/22195 A2**

- (51) International Patent Classification<sup>7</sup>: **G06F**
- (21) International Application Number: **PCT/US00/26097**
- (22) International Filing Date:  
22 September 2000 (22.09.2000)
- (25) Filing Language: **English**
- (26) Publication Language: **English**
- (30) Priority Data:  
09/405,608 24 September 1999 (24.09.1999) **US**
- (71) Applicant: **AKAMBA CORPORATION [US/US];**  
15595 Los Gatos Blvd., Los Gatos, CA 95032 (US).
- (72) Inventors: **SMITH, Jack, J.**; 28021 Audry Smith Lane,  
Saratoga, CA 95070 (US). **BURRIGHT, Richard, T.**;  
2258 Dry Creek Ct., San Jose, CA 95124 (US). **WOR-**  
**LEY, W., Spencer, III**; 311 Correas Avenue, Half Moon

Bay, CA 94019 (US). **MACDONELL, Eoin, B.**; 1057  
Yorktown Dr., Sunnyvale, CA 94087 (US). **VASTANO,**  
**John, A.**; 3497 Cabrillo Ave., Santa Clara, CA 95051  
(US). **WEATHERFORD, William, T.**; 1556 Forge Rd.,  
San Mateo, CA 94402 (US).

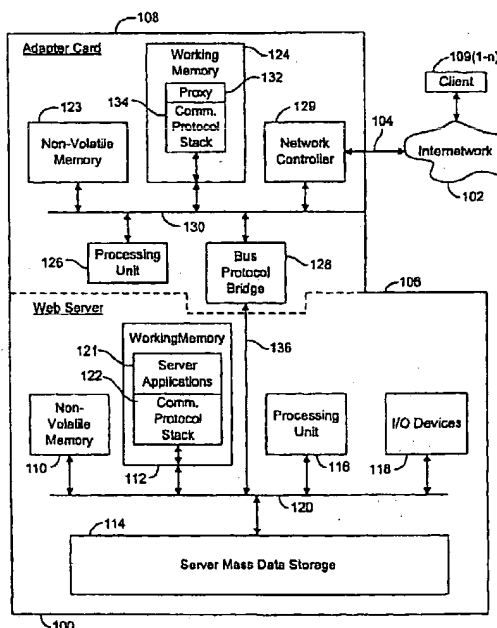
(74) Agent: **HENNEMAN, Larry, E., Jr.**; Henneman & Saun-  
ders, 121 E. 11th Street, Tracy, CA 95376 (US).

(81) Designated States (*national*): AE, AL, AM, AT, AU, AZ,  
BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK,  
DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL,  
IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU,  
LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT,  
RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA,  
UG, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM,  
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian  
patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European

[Continued on next page]

(54) Title: **SYSTEM AND METHOD FOR MANAGING CONNECTIONS BETWEEN A CLIENT AND A SERVER**



(57) Abstract: An adapter card (108) for managing connections between clients (109) and a network server (106) off-loads the connection management burden from the server (106). The adapter card (108) includes a memory (124) with an embedded proxy application (132) and a communication protocol stack (134), a processing unit (126) for executing the application code, a network controller (129) for interfacing with an internetwork (102), and a bus protocol bridge (128) for interfacing with the internal bus (120) of the network server (106). The proxy application (132) receives client requests on behalf of the server (106) over relatively slow and unreliable network connections, and submits the requests to the server over fast, reliable bus connections.



patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

**Published:**

— *Without international search report and to be republished upon receipt of that report.*

## SYSTEM AND METHOD FOR MANAGING CONNECTIONS BETWEEN A CLIENT AND A SERVER

5 Jack J. Smith; Richard T. Burrig; W. Spencer Worley, III;  
Eoin B. MacDonell; John A. Vastano; and William T. Weatherford

### BACKGROUND OF THE INVENTION

#### Field of the Invention

10 This invention relates generally to network servers and more particularly to servers that host a large number of client connections. Even more particularly, the present invention relates to servers (e.g., internet web servers) which host a large number of relatively slow client connections.

#### Description of the Background Art

15 It is common for network file servers such as internet web servers to host a large number of relatively slow client connections. The large number of open connections places a substantial burden on the server central processing unit (CPU), just to manage the open connections. For example, managing the open connections on a loaded server can consume 30-40 % or more of the CPU's operating capacity. This burden substantially  
20 reduces the percentage of CPU cycles available to perform the primary function of the server, i.e., providing data to clients.

The connection management burden on the server CPU degrades the performance of the server software routines and reduces the maximum number of client connections that can be open at one time. As a result, web-hosting companies must provide additional, redundant  
25 servers to serve an increased number of clients. The cost of acquiring and maintaining additional web servers is substantial.

Proxy servers perform some client connection management functions, and are known in the art. However, it is well-known and commonly accepted in the art that such proxy servers must be housed separately from the server, and thus must communicate with the  
30 server over relatively slow, error prone network connections which the server must manage.

See for example, Ari Luotonen, *Web Proxy Servers* (Prentice Hall, 1997), which is incorporated herein by reference.

What is needed, therefore, is a system and method for relieving the server CPU of the connection management burden, thus allowing the server to more efficiently host an increased  
5 number of clients.

### SUMMARY

The present invention overcomes the problems associated with the prior art by providing a system and method for managing connections between a plurality of clients and a  
10 server. The invention facilitates off-loading the connection management burden from the host CPU to an adapter card interposed between the network and the host bus.

The adapter card includes a network controller, a memory device, a processing unit, and a protocol adapter. The memory device provides storage for data and code. The code includes a proxy application that communicates with clients on the network via the network  
15 controller, and communicates with the server via the protocol adapter, which is coupled directly to the server bus.

When executed by the processing unit, the proxy application manages client connections by establishing network connections between the proxy application and clients via the network, and by establishing bus connections between the proxy application and the  
20 server via the server bus. Additionally, the memory device provides data buffering, which allows many network connections to be open with clients, while a relatively few bus connections are open to the server. In a particular embodiment, the proxy accumulates client data in the buffers from the large number of slow client connections, and then submits the client data to the server over the fast bus connections. Conversely, the proxy receives server  
25 data via the fast bus connections, temporarily stores the server data, and then forwards the server data to the clients via the slow client connections.

In a more particular embodiment, the code includes a communications protocol stack that is employed by the application proxy to communicate with the clients and the server. In an even more particular embodiment, the communications protocol stack is a Transmission  
30 Control Protocol/Internet Protocol (TCP/IP) stack.

In one embodiment, the server connections are opened only after the proxy determines that a complete client request has been received. The server connections are then closed after the proxy receives a response to the client request from the server. Optionally, a predetermined number of persistent server connections are opened at system start-up, and the proxy uses these persistent connections to communicate with the server.

The proxy application optionally includes a number of application specific proxies, including but not limited to an HTTP proxy, a security proxy, and/or a pass-through proxy. In a particular embodiment, a master process module of the proxy discerns an application identifier (e.g., a well known port number) from the client data, and invokes one or more of the application specific proxies corresponding to the value of the identifier.

### BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is described with reference to the following drawings, wherein like reference numbers denote substantially similar elements:

FIG. 1 is a block diagram of a server and an adapter card according to the present invention;

FIG. 2 is a block diagram of the working memory of the adapter card of FIG. 1, showing the proxy module in greater detail;

FIG. 3 is a block diagram showing the application proxies module of FIG. 2 in greater detail;

FIG. 4 is a block diagram showing exemplary data structures for at least some of the data stored in the data buffer of FIG. 2;

FIG. 5 is a flow chart summarizing one method for managing connections between a client and a server according to the present invention;

FIG. 6 is a flow chart summarizing one method for performing the first step of the method of FIG. 5;

FIG. 7 is a flow chart summarizing one method for performing the second step of the method of FIG. 5;

FIG. 8 is a flow chart summarizing one method for performing the third step of the method of FIG. 5;

FIG. 9 is a flow chart summarizing one method for performing the fourth step of the method of FIG. 5;

FIG. 10 is a flow chart summarizing one method for performing the fifth step of the method of FIG. 5; and

5        FIG. 11 is a flow chart summarizing one method for performing the sixth step of the method of FIG. 5.

### DETAILED DESCRIPTION

The present invention overcomes the problems associated with the prior art, by off-loading much of the connection management burden from the server's main processor with a proxy application run on a different processing unit. In the following description, numerous specific details are set forth (e.g., particular communications protocols, particular software and data structures, etc.) in order to provide a thorough understanding of the invention. Those skilled in the art will recognize, however, that the invention may be practiced apart from these specific details. In other instances, details of well known network components and programming practices (e.g., establishing connections via a communications protocol stack) have been omitted, so as not to unnecessarily obscure the present invention.

FIG. 1 is a block diagram showing a system 100 coupled to an internetwork 102 via a physical network media 104. In a particular implementation, system 100 is an internet web server, and internetwork 102 is the Internet, but those skilled in the art will recognize that the present invention may be implemented in any type of network server.

System 100 includes a file server (e.g., an HTTP web server) 106 and an adapter card 108. File server 106 provides data to and receives data from clients 109(1-n) on internetwork 102, via adapter card 108. Adapter card 108 establishes and maintains network connections between clients 109(1-n) and adapter card 108, and establishes bus connections between server 106 and adapter card 108. Thus connected, adapter card 108 receives communications from clients 109(1-n) on behalf of server 106, forwards the communications to server 106, receives responses from server 106 on behalf of clients 109, and forwards the responses to clients 109.

30        Server 106 includes non-volatile memory 110, working memory 112, server mass data storage 114, a processing unit 116, and one or more user input/output (I/O) devices 118, all

intercommunicating via a server bus 120 (e.g., PCI bus). Non-volatile memory 110 (e.g., read-only memory and/or one or more hard-disk drives) provides storage for data and code which is retained even when server 106 is powered down. Working memory 112 (e.g., random access memory) provides operational memory for server 106, and includes executable code (e.g., an operating system) which is loaded into working memory 112 during start-up. Among other programs, working memory 112 includes server applications 121 and a communication protocol stack 122. Server applications 121 include network software applications (e.g., FTP, HTTP, etc.) which allow server 106 to function as a network server. Communications protocol stack 122 is a standard protocol stack (e.g., TCP/IP) which facilitates communication with other machines over an internetwork. Standard protocol stacks are well known in the art. See, for example, W. Richard Stevens, *TCP/IP Illustrated, Vol. 1* (Addison-Wesley, 1994), which is incorporated herein by reference. Server mass data storage 114 provides data storage (e.g., one or more hard disk drives) for data (e.g., HTML pages, graphics files, etc.), which the server provides to clients 109(1-n) attached to internetwork 102. Processing unit 116 executes the instructions in working memory 112 to cause server 106 to carry out its primary function (e.g., providing data to and receiving data from clients). I/O devices 118 typically include a keyboard, a monitor, and/or such other devices which facilitate user interaction with server 106. Each of the above described components is typically found in a network server such as an internet web server.

Adapter card 108 includes non-volatile memory 123, working memory 124, a processing unit 126, a bus protocol bridge 128, and a network controller 129, all intercommunicating via an adapter bus 130. Non-volatile memory 123 provides storage for data and code (e.g., boot code) which is retained even when adapter 108 is powered down. Processing unit 126 imparts functionality to adapter card 108 by executing the code present in working memory 124. Bus protocol bridge 128 provides an interface between adapter bus 130 and server bus 120, and network controller 129 provides an interface between adapter bus 130 and network media 104.

Working memory 124 provides operational memory for adapter 108, and includes a proxy application 132 and a communication protocol stack 134. Proxy 132 and protocol stack 134 are loaded from non-volatile memory 123 into working memory 124 at start-up. Optionally, proxy 132 and protocol stack 134 can be loaded from one or more alternative



sources, including but not limited to non-volatile memory 110 or server mass data storage 114 of server 106. Proxy 132, when executed by processing unit 126, establishes and manages the above described connections between adapter 108 and server 106 and between adapter 108 and clients 109.

5        In this particular embodiment of the invention, protocol stacks 122 and 134 are standard (e.g., TCP/IP) protocol stacks. Employing a standard communication protocol stack in adapter 108 facilitates the use of the standard communication software (e.g., protocol stack 122) already present in the vast majority of network servers. Those skilled in the art will recognize, however, that this particular element (as well as other described elements, even if  
10   not explicitly stated) is not an essential element of the present invention. For example, the present invention may be practiced with custom communication software (e.g., direct communication between server applications 121 and either protocol stack 134 or proxy 132) in both server 106 and adapter 108. Further, in particular embodiments of the invention, this element may be omitted by providing proxy 132 with direct access to the resources (e.g.,  
15   server mass data storage 114) of server 106.

Adapter card 108 is coupled to server 106 via a bus connection 136 between bus protocol bridge 126 and server bus 120. In this particular embodiment, bus connection 136 is a typical bus expansion slot, for example a PCI slot. Those skilled in the art will recognize, however, that the present invention may be implemented with other types of bus connections,  
20   including but not limited to an ISA slot, a USB port, a serial port, or a parallel port. Bus connection 136 facilitates high speed, large packet size, relatively error free (as compared to network connections) communication between proxy 132 and server applications 121, greatly reducing the connection management burden on processing unit 116 of server 106. In summary, proxy 132 (running on processing unit 116) communicates with clients 109 over  
25   slow, error prone network connections, and then communicates with server applications 121 on behalf of clients 109 over high speed bus connection 136.

FIG. 2 is a block diagram of working memory 124 showing proxy 132 and protocol stack 134 in greater detail. Those skilled in the art will recognize that while the various software modules of proxy 132 are shown as interconnected functional blocks, the software  
30   modules are actually blocks of executable code stored in working memory 124 that can communicate with one another when executed by processing unit 126 (FIG. 1).

Proxy 132 includes a master process module 202, a plurality of client process modules 204(1-n), a data buffer 206, and an application proxies module 208. Master process module provides overall control and coordination of the various modules of proxy 132. Responsive to a connection request from a client 109 on internetwork 102 (FIG. 1) master process 202

5 accepts the connection request, initializes a data structure for that client connection in data buffer 206, initiates a new, separate client connection process 204 to handle the connection, and then notifies application proxies 208 that the particular client connection has been established. Each client process 204 handles one such client connection. Application proxies 208 establish and manage bus connections with server 106. Data buffer 206 provides storage

10 for data received from clients 109 and destined for server 106, for data received from server 106 and destined for clients 109, and for connection data relating to established client and/or server connections.

Communications protocol stack 134 is a TCP/IP stack including a sockets layer 210, a TCP layer 212, an IP layer 214, and a device layer including a network driver 216 and a

15 server bus driver 218. The functionality of each of the individual layers of protocol stack 134 is well known in the art, and will not, therefore, be discussed in detail herein. Connections between the various modules of proxy 132 and server applications 121 are established through sockets layer 210, TCP layer 212, IP layer 214 and server bus driver 218. Connections between the various modules of proxy 132 are established with clients 109

20 through sockets layer 210, TCP layer 212, IP layer 214 and network driver 216.

FIG. 3 is a block diagram showing application proxies module 208 to include a plurality of application specific proxies 208(1-f), including a hypertext transfer protocol (HTTP) proxy 208(1), a pass-through proxy 208(2), a security proxy 208(3), and an "other" proxy 208(f). Master process 202 notifies application proxies 208 of an established client

25 connection, by configuring one or more of the application specific proxies 208(1-f) to service the client connection. One means of configuring an application specific proxy (e.g., HTTP proxy 208(1)) is to enter a client process identifier in the processing queue of the application specific proxy.

Master process 202 determines which of the application specific proxies to implement

30 for a particular client process from the port number included in the client connection request. It is standard practice to use well known port numbers to identify particular network

applications and/or protocols (e.g., file transfer protocol (FTP), HTTP, etc.). For example, port number 80 corresponds to an HTTP connection request. Master process 202 therefore notifies HTTP proxy 208(1) of all client process' 204 initiated in response to a connection request indicating port 80.

5 HTTP proxy 208(1) monitors each of the client processes of which it is notified. When HTTP proxy 208(1) determines that a complete HTTP request is received and stored in data buffer 206 by a client process (e.g., 204(n)), HTTP proxy 208(1) opens a connection to the server, transmits the request to the server, receives a response from the server, stores the response in data buffer 206 and then closes the server connection. The server response is then  
10 transmitted to client 109(n) by the associated client process 204(n).

When master process 202 receives a connection request with a port number that does not correspond to any of the other application specific proxies, master process 202 notifies pass-through proxy 208(2). Pass-through proxy 208(2) simply opens a server connection, transfers the data received from the associated client process 204 from data buffer 206 to  
15 server 106, and then closes the server connection.

Master process 202 may notify some application specific proxies of all client connections, regardless of the associated port number. For example, security proxy 208(3) is operative to screen all client connection requests by, for example, terminating any client process initiated in response to a connection request lacking some indicia of authorization,  
20 prior to implementing one of the other application specific proxies.

"Other" proxy 208(f) is included in FIG. 3 to show that application proxies 208 can include any currently known or future developed proxy that is desirable for a particular application, including but not limited to, caching HTTP proxy applications, electronic mail applications, and file transfer applications.

25 FIG. 4 shows an example of client data structures 402(1-n) and proxy data structures 404(1-f), implemented in data buffer 206 to effect the transfer of data through proxy 132. Master process 202 creates and initializes one client data structure 402 for each client process 204, and one proxy data structure 404 for each application specific proxy in application proxies 208.

30 Each client data structure 402 includes a client socket 406, a server socket 408, a connection state 410, an input queue 412, an output queue 414, and application proxy data

416. For each client connection (e.g., connection (n)), client socket 406(n) and server socket 408(n) each include the IP address and port number of the client 109(n) and server 106, respectively, thus uniquely associating each client data structure 402(n) with a single one of client processes 204(n). Connection state 410(n) indicates the current status (e.g., complete request received, response received, etc.) of the connection (n). Input queue 412(n) is used to store and accumulate data received from client 109(n) by the client process 204(n) associated with the particular data structure 402(n). Output queue 414(n) is used to store data from application proxies 208 which is to be forwarded to client 109(n) by client process 204(n). Application proxy data 416(n) is provided to store any information specific to a particular application proxy (e.g., flags, etc.).

Each proxy data structure (e.g., 404(f)) includes a client queue 418(f), a client ready queue 420(f), and a read pending queue 422(f). Client queue 418(f) includes a client process descriptor (e.g., a pointer to a related client data structure 402) for each client process 204 associated with the particular application proxy (f) to which the proxy data structure 404(f) corresponds. Client ready queue 420(f) includes a client process descriptor for each client data structure 402 that has data in its input queue 412 that is ready to be processed (e.g., transferred to server 106) by the associated application proxy (f). Read pending queue 422(f) includes the client process descriptor for each client process that is awaiting a response from server 106.

Those skilled in the art will understand that the above described client data structure 402 and proxy data structure 404 are exemplary in nature, and that other data structures may be employed with the present invention. The configuration of such alternate data structures will necessarily depend on the function and structure of the particular application specific proxies that are employed.

FIG. 5 is a flowchart summarizing a particular method 500 of managing connections between clients and a server according to the present invention. In a first step 502 proxy 132 establishes a network connection with a client 109, and then in a second step 504 receives a communication (e.g., an HTTP request) from client 109 via the network connection. Next, in a third step 506, proxy 132 establishes a bus connection with server 106, and then in a fourth step 508 forwards the received client communication to server 106 via the bus connection. Then, in a fifth step 510, proxy 132 receives a response (e.g., HTML data) to the client

communication from server 106, and in sixth step 512 transmits the response to client 109 via the client network connection. Finally, in a seventh step 514, proxy 132 determines whether there is a signal to terminate (e.g., shut down), and if there is a signal to terminate, then method 500 ends. If there is no signal to terminate in seventh step 514, then method 500  
5 returns to first step 502, to establish a network connection with another client 109.

FIG. 6 is a flow chart summarizing one particular method 600 of performing first step 502 (establishing a network connection with a client) of method 500. In a first step, master process 202 connects to internetwork 102. Then in a second step 604 master process 202 listens to the traffic on internetwork 102 to determine if there is a connection request from a  
10 client 109. If there is no client connection request, then method 600 ends. If there is a connection request from a client 109(n), then in a third step 606, master process 202 accepts the connection request from client 109(n), initiates a client process 204(n) to handle the connection, and initializes a client data structure 402(n) in data buffer 206. Next, in a fourth  
15 step 608, master process 202 discerns a proxy application identifier (e.g., a port number) from the client connection request and notifies one or more of application proxies 208(1-f), depending on the value of the identifier, by writing the client process descriptor (e.g., pointer to client data structure 402(n)) into the client queues 418 of the respective proxy data  
structures 404. Finally, in a fifth step 610, master process 202 determines whether the maximum allowed number of client connections are open. If the maximum number of client  
20 connections are open, then method 600 ends. If the maximum number of client connections are not open, then method 600 returns to second step 604, to listen for another connection request.

FIG. 7 is a flow chart summarizing a method 700 of performing second step 504 (receiving a communication from a client 109) of method 500. In a first step 702, master  
25 process 202 determines whether there are any client processes 204 to be processed to receive data. If master process 202 has already processed all of client processes 204(1-n), then method 700 ends. If not, then in a second step 704, master process 202 calls the first client process 204(1). Then, in a third step 706, client process 204(1) checks its client connection (e.g., the TCP buffer) to determine whether there is any data coming in from client 109(1). If  
30 there is no incoming data for the first client process 204(1), then method 700 returns to first step 702 to process any remaining client processes 204(2-n). If, in third step 706, client

process 204(1) determines that there is incoming data from client 109(1), then in a fourth step 708, client process 204(1) checks client data structure 402(1) to determine if input queue 412(1) is available to receive client data. If input queue 412(1) is not available, then method 700 returns to first step 702 to process any remaining client processes 204(2-n). If in fourth  
5 step 708, client process 204(1) determines that input queue 412(1) is available to receive data, then in a fifth step 710, client process 204(1) transfers the incoming client data into input queue 412(1). Then, in a sixth step 712, client process 204(1) determines whether the data accumulated in input queue constitutes a complete request (i.e., data ready to be transferred to server 106, for example a complete HTTP request). If the data does not constitute a complete  
10 request, then method 700 returns to first step 702 to process any remaining client processes 204(2-n). If, however, client process 204(1) determines in sixth step 712 that the data in input queue 412(1) constitutes a complete request, then, in a seventh step 714, client process notifies proxy applications 208 that there is a complete request by, for example, setting connection state 410(1) to so indicate. Then, method 700 returns to first step 702 to  
15 determine whether there are any more client processes 204(2-n) to process.

FIG. 8 is a flow chart summarizing a method 800 of performing third step 506 (establishing a bus connection with server 106) of method 500. In a first step 802, a first one of application proxies 208(1) retrieves the first client descriptor from client queue 418(1) of its proxy data structure 404(1). Then in a second step 804, proxy 208(1) checks the  
20 connection state 412 of the client data structure 402 identified by the first client descriptor to determine if the first client has a complete request in its input queue 412. If connection state 412 indicates a complete request, then in a third step 806, proxy 208(1) adds the client descriptor to its client ready queue 420(1). Next, in a fourth step 808, proxy 208(1) determines whether the maximum number of connections to server 106 are open. If the  
25 maximum number of server connections are already open, then method 800 ends. If the maximum number of server connections are not already opened, then in a fifth step 810, proxy 208(1) opens a bus connection with server 106 and writes connection information in server socket 408 of the associated client data structure 402. Next, in a sixth step 812, proxy 208(1) determines whether it has checked the last client descriptor in its client queue 418(1).  
30 If the last descriptor has been checked, then method 800 ends. Otherwise, method 800 returns to first step 802 to retrieve the next descriptor in client queue 418(1). If, in second

step 804, proxy 208(1) determines that a complete client request has not been received, then method 800 proceeds directly to sixth step 812. Once all of the descriptors in client queue 418(1) of proxy data structure 404(1) have been processed, method 800, or a similar method, is repeated for each of the other application proxies 208(2-f).

5        FIG. 9 is a flow chart summarizing a method 900 of performing fourth step 508 (forwarding a client communication to server 106) of method 500. In a first step 902, proxy 208(1) retrieves the first client descriptor from the client ready queue 420(1) of its proxy data structure 404(1). Then, in a second step 904, proxy 208(1) checks the server socket 408 of the first client's data structure 402 to determine whether a server connection is open. If a  
10    server connection is open, then in a third step 906, proxy 208(1) transfers the client data (e.g., HTTP request) from the client input queue 412 to server 106 over the open server connection. Next, in a fourth step 908, proxy 208(1) moves the client descriptor from the client ready queue 420(1) to the read pending queue 422(1). Then in a fifth step 910, proxy 208(1) determines whether the last client in the client ready queue 420(1) has been checked. If not,  
15    then method 900 returns to first step 902 to check the next client in client ready queue 420(1). If the last client has been checked, then method 900 ends. If, in second step 904, proxy 208(1) determines that there is no server connection open for a particular client, then method 900 proceeds directly to fifth step 910 to determine whether the last client in the client ready queue 420(1) has been checked. Once all of the descriptors in client ready queue 420(1) of  
20    proxy data structure 404(1) have been processed, method 900, or a similar method, is repeated for each of the other application proxies 208(2-f).

      FIG. 10 is a flow chart summarizing a method 1000 of performing fifth step 510 (receive a response from server 106) of method 500. In a first step 1002, proxy 208(1) determines whether read pending queue 422(1) is empty, and if so method 1000 ends. If read  
25    pending queue 422(1) is not empty, then in a second step 1004 proxy 208(1) retrieves the first client descriptor from the read pending queue. Next, in a third step 1006, proxy 208(1) checks the open server connection identified in server socket 408 of the client data structure 402 identified by the first client descriptor, to determine whether there is any incoming server data (i.e., a response to the client request) on that connection. If there is no incoming server  
30    data on that connection, then method 1000 returns to second step 1004 to check the next client in the read pending queue. If there is incoming server data, then in a fourth step 1008,

proxy 208(1) checks to determine whether output queue 414 of the client data structure 402 identified by the client descriptor is available. If output queue 414 is not available, then method 1000 returns to second step 1004 to check the next client descriptor in read pending queue 422(1). If output queue 414 is available, then, in a fifth step 1010, proxy 208(1) moves  
5 the incoming server data into the output queue 414 of client data structure 402. Next, in a sixth step 1012, proxy 208(1) determines whether the server data includes an "end of file" indicator. If not, then in a seventh step 1014, proxy 208(1) checks to determine whether the last client descriptor in read pending queue 422(1) has been processed. If so, method 1000 ends. If not, method 1000 returns to step 1004 to read the next client descriptor from read  
10 pending queue 422(1).

If, in sixth step 1012, proxy 208(1) determines that the server data includes an end-of-file indicator, then method 1000 proceeds to an eighth step 1016, wherein proxy 208(1) removes the client descriptor from the read pending queue, and then in a ninth step 1018 closes the server connection. After ninth step 1018, method 1000 returns to seventh step  
15 1014. Once all of the descriptors in read pending queue 422(1) of proxy data structure 404(1) have been processed, method 1000, or a similar method, is repeated for each of the other application proxies 208(2-f).

FIG. 11 is a flow chart summarizing a method 1100 of performing sixth step 512 (transmitting data to clients 109) of method 500. In a first step 1102, master process 202 calls  
20 first client process 204(1). Then, in a second step 1104, first client process 204(1) determines whether there is any data in output queue 414(1) of client data structure 402(1). If there is no data in output queue 414(1), then method 1100 returns to first step 1102 where master process 202 calls the next of the remaining client processes 204(2-n). If, however, in second step 1104, client process 204(1) determines that there is data in output queue 414(1), then in a  
25 third step 1106, client process 204(1) determines whether the client network connection is ready to receive data. If the client network connection is not ready, then method 1100 returns to first step 1102. If the client network connection is ready, then in a fourth step 1108, client process 204(1) moves at least a portion of the data in output queue 414(1) to the client connection (e.g., the TCP output buffer). Next, in a fifth step 1110, master process 202  
30 determines whether the last client process has been called. If so, method 1100 ends. If not,



method 1100 returns to first step 1102 to call the next of the remaining client processes 203(2-n).

The description of particular embodiments of the present invention is now complete. Many of the described features may be substituted, altered or omitted without departing from the scope of the invention. For example, the operative components of adapter 108 (e.g., processing unit 126 and proxy 132) can be incorporated directly into a server instead of being provided in a removable adapter card. Further, alternate data structures may be substituted for the exemplary data structures provided. Additionally, the particular orders of methods and routines disclosed herein are not considered to be essential elements of the present invention. As yet another example, master process 202 can be configured to open a predetermined number of persistent bus connections with server 106 at start-up, and manage the use of those connections by application proxies 208(1-f), thus eliminating the need for server 106 to repetitively open and close the bus connections. These and other deviations from the particular embodiments shown will be apparent to those skilled in the art, particularly in view of the foregoing disclosure.

We claim:

1. A method for managing connections between at least one client and a server, said method comprising:
  - 5 establishing a network connection with one of said clients via a network ;
  - receiving a communication from said client via said network connection;
  - establishing a bus connection with said server via an internal bus of said server; and
  - forwarding said client communication to said server via said bus connection.
- 10 2. A method according to Claim 1, wherein said step of receiving a communication from said client includes storing said communication in a buffer.
3. A method according to Claim 2, wherein said step of storing said communication in a buffer includes accumulating one or more separate transmissions from said client in said  
15 buffer.
4. A method according to Claim 3, wherein said step of establishing a bus connection with said server includes waiting until a complete client request is accumulated in said buffer before establishing said bus connection with said server.  
20
5. A method according to Claim 4, further comprising:
  - receiving a response to said client communication from said server via said bus
  - connection; and
  - forwarding said response to said client via said network connection.  
25
6. A method according to Claim 5, wherein said step of receiving said response from said server includes storing said response in a buffer.
7. A method according to Claim 6, wherein said step of receiving said response from  
30 said server includes terminating said bus connection after said response is received.

8. A method according to Claim 1, further comprising:  
receiving a response to said client communication from said server via said bus  
connection; and  
forwarding said response to said client via said network connection.

5

9. A method according to Claim 8, wherein said step of receiving said response from  
said server includes storing said response in a buffer.

10. A method according to Claim 9, wherein said step of receiving said response from  
said server includes terminating said bus connection after said response is received.

11. A method according to Claim 8, wherein said client communication includes an  
HTTP request.

12. A method according to Claim 11, wherein said response from said server includes  
an HTML page.

13. A method according to Claim 1, wherein said step of establishing a network  
connection with a client includes establishing a separate network connection with each of a  
plurality of clients via said network.

14. A method according to Claim 13, wherein said step of establishing said bus  
connection with said server includes establishing a plurality of connections with said server  
via said internal bus of said server.

25

15. A method according to Claim 14, wherein the maximum number of simultaneous  
client connections exceeds the maximum number of simultaneous server connections.

16. A method according to Claim 1, further comprising performing a security  
operation on said client communication prior to forwarding said client communication to said  
server.

30

17. A method according to Claim 1, wherein:

said step of receiving said client communication includes discerning an application  
identifier from said client communication; and

5 said step of forwarding said client communication to said server includes invoking  
one of a plurality of proxy applications based on said application identifier.

18. A method according to Claim 17, wherein said application identifier is the  
connection port number.

10

19. A method according to Claim 1, wherein said step of receiving said client  
communication includes receiving at least a portion of an HTTP request.

20. A computer readable medium having code embodied therein for causing an  
15 electronic device to perform the steps of Claim 1.

21. A computer readable medium having code embodied therein for causing an  
electronic device to perform the steps of Claim 2.

20 22. A computer readable medium having code embodied therein for causing an  
electronic device to perform the steps of Claim 3.

23. A computer readable medium having code embodied therein for causing an  
electronic device to perform the steps of Claim 4.

25

24. A computer readable medium having code embodied therein for causing an  
electronic device to perform the steps of Claim 5.

25. A computer readable medium having code embodied therein for causing an  
30 electronic device to perform the steps of Claim 6.

26. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 7.

27. A computer readable medium having code embodied therein for causing an  
5 electronic device to perform the steps of Claim 8.

28. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 9.

10 29. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 10.

30. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 11.

15

31. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 12.

32. A computer readable medium having code embodied therein for causing an  
20 electronic device to perform the steps of Claim 13.

33. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 14.

25 34. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 15.

35. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 16.

30

36. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 17.

37. A computer readable medium having code embodied therein for causing an  
5 electronic device to perform the steps of Claim 18.

38. A computer readable medium having code embodied therein for causing an electronic device to perform the steps of Claim 19.

10 39. An adapter card for coupling a server with an internal bus to a network, said adapter card comprising:  
a network controller for communicating with clients on said network;  
a memory device for storing data and code, said code including a proxy application;  
a processing unit coupled to said memory device for executing said code; and  
15 a protocol adapter coupled to said processing unit, and adapted to couple to said internal bus, for communicating with said server.

40. An adapter card according to Claim 39, wherein said code further comprises a communication protocol stack.

20 41. An adapter according to Claim 40, wherein said communication protocol stack comprises a standard TCP/IP protocol stack.

42. An adapter card according to Claim 39, wherein said proxy application includes a  
25 security proxy.

43. An adapter card according to Claim 39, wherein said proxy application includes a pass-through proxy.

30 44. An adapter card according to Claim 39, wherein said proxy application includes an HTTP proxy.

45. An adapter card according to Claim 39, further comprising a data buffer for storing data received from said clients.

5           46. An adapter card according to Claim 39, wherein said proxy application includes a master process module responsive to a connection request received from one of said clients, and operative to establish a connection with said client and to initiate a new client process module to maintain said established connection.

10           47. An adapter card according to Claim 46, wherein said master process module is further operative to notify said proxy application of said established connection.

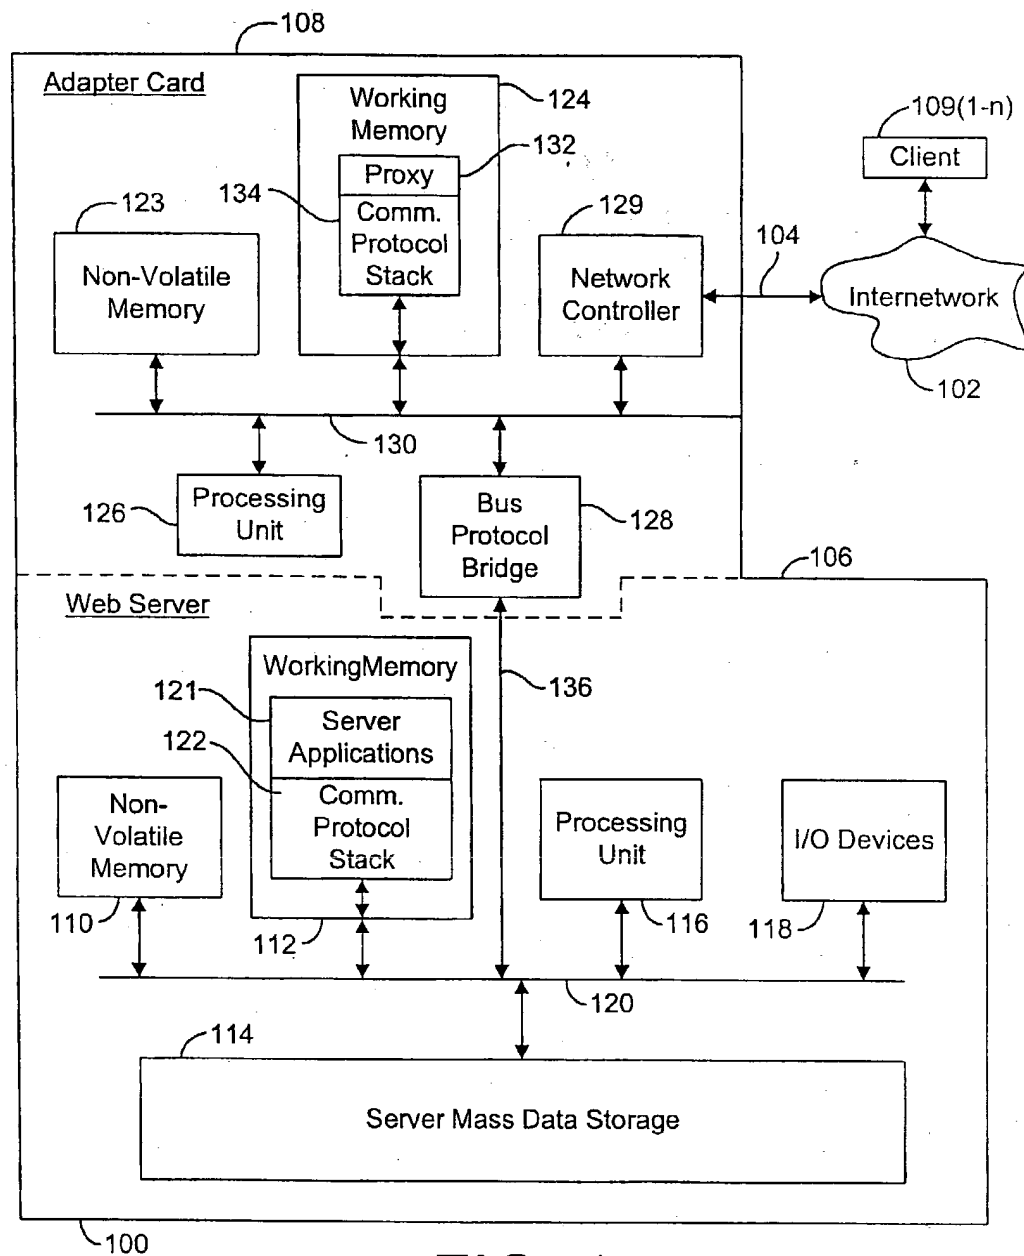


FIG. 1



FIG. 2

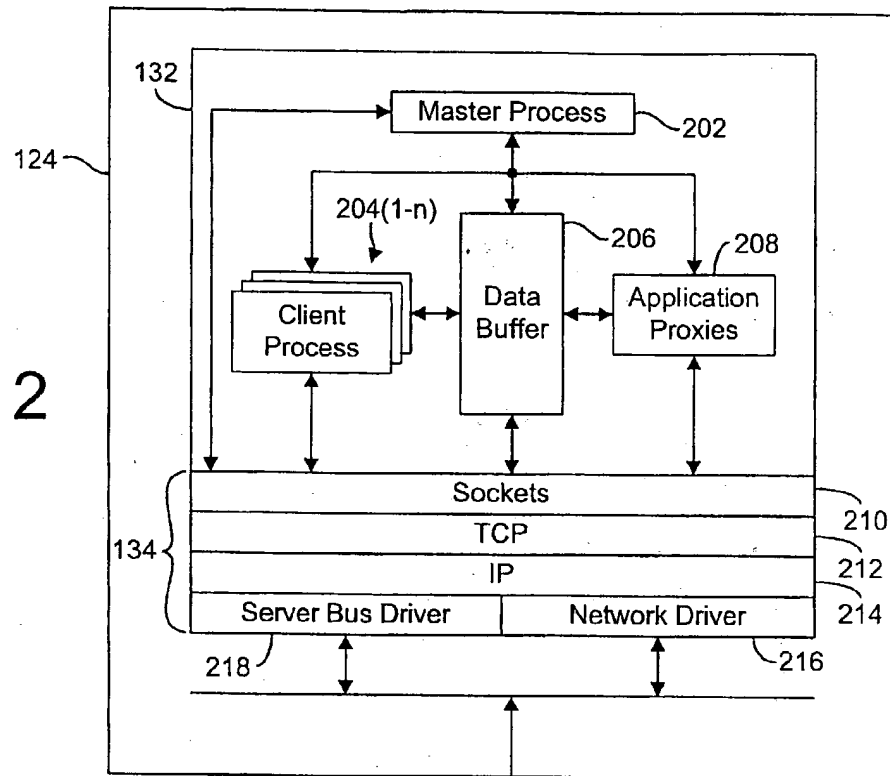


FIG. 3


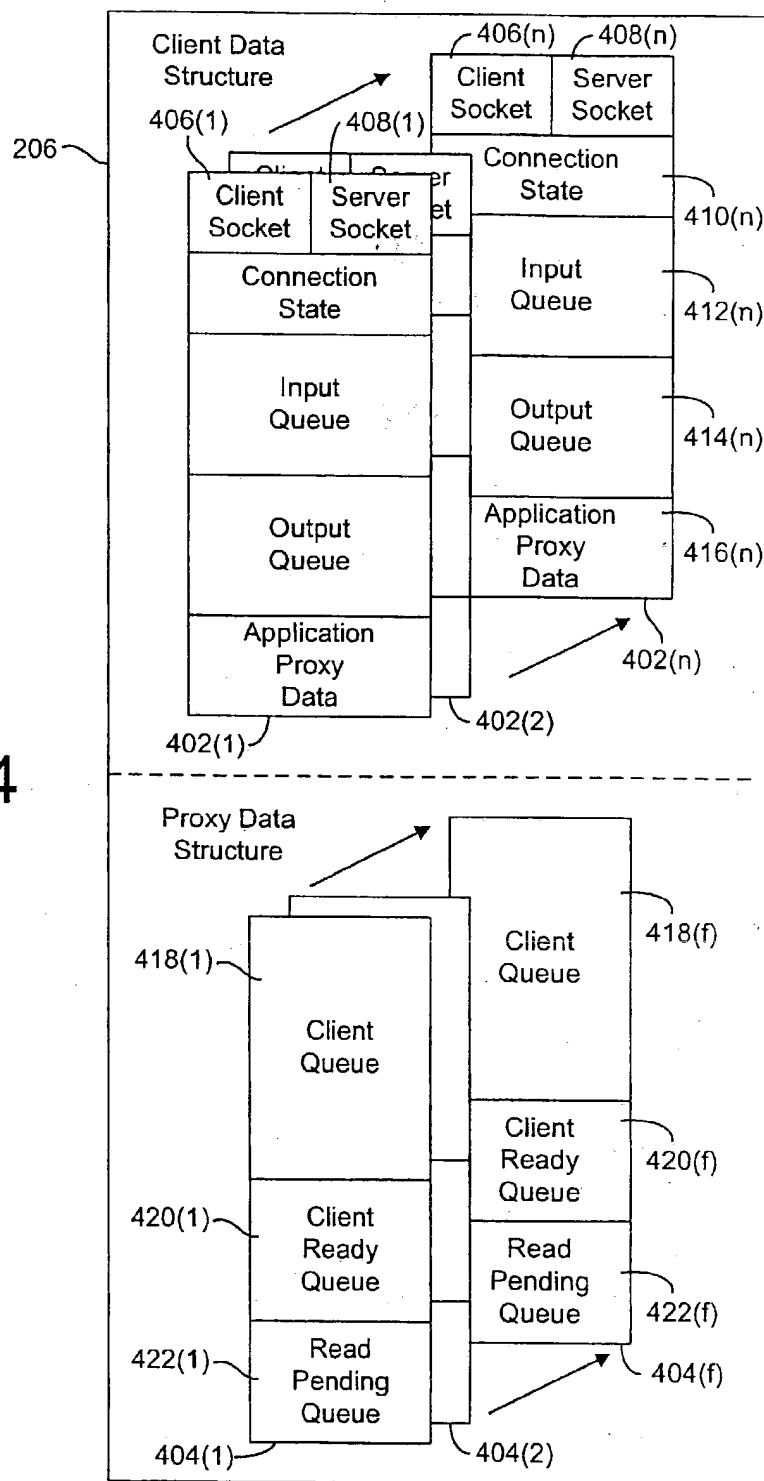


FIG. 4



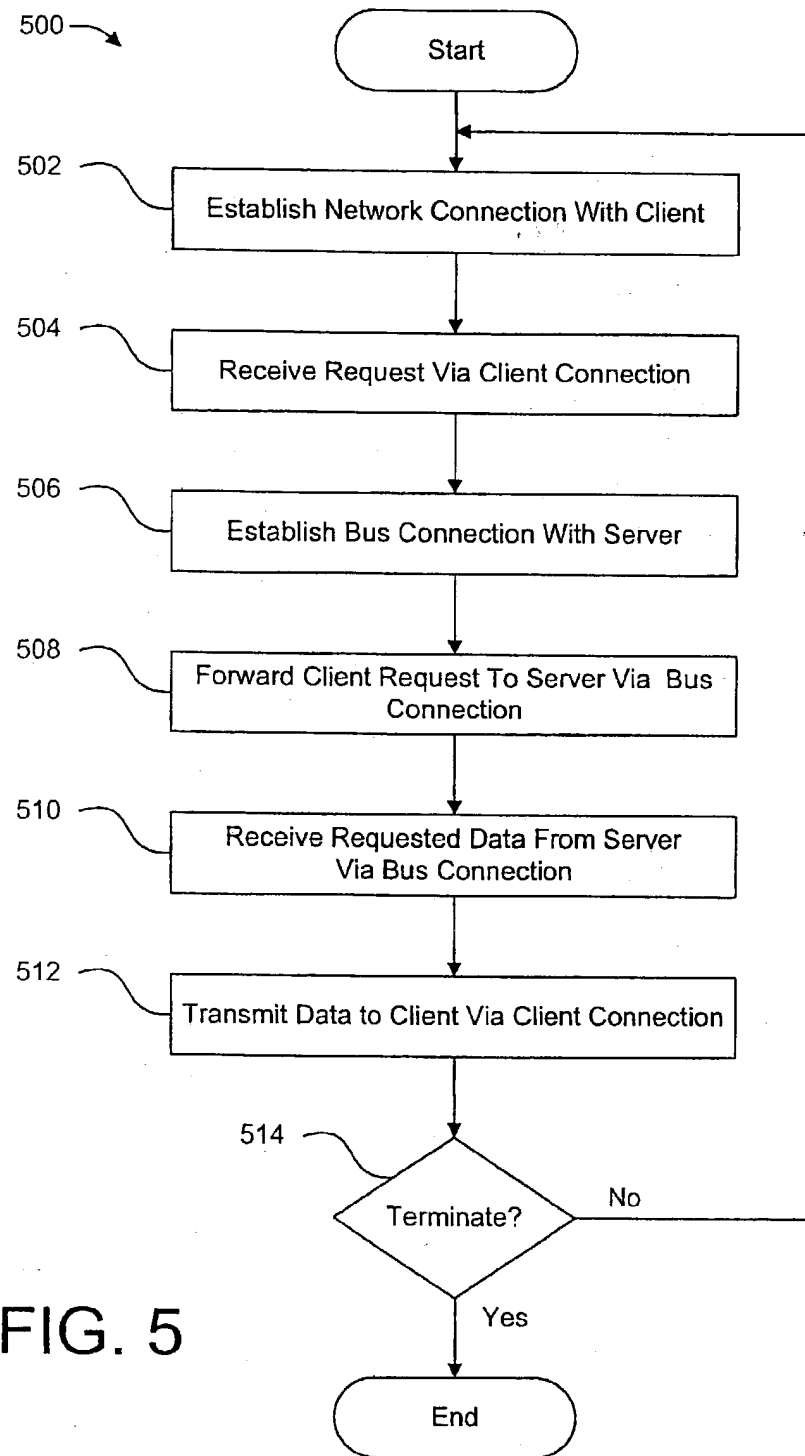


FIG. 5

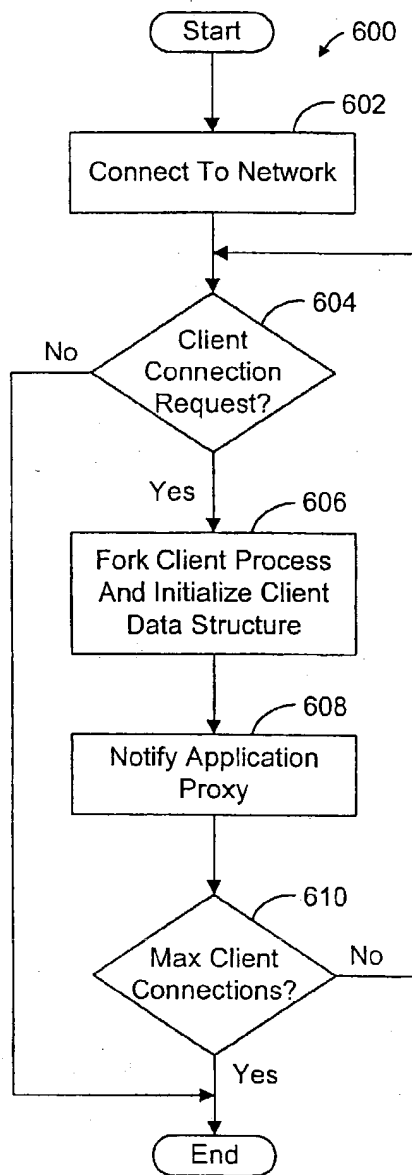


FIG. 6

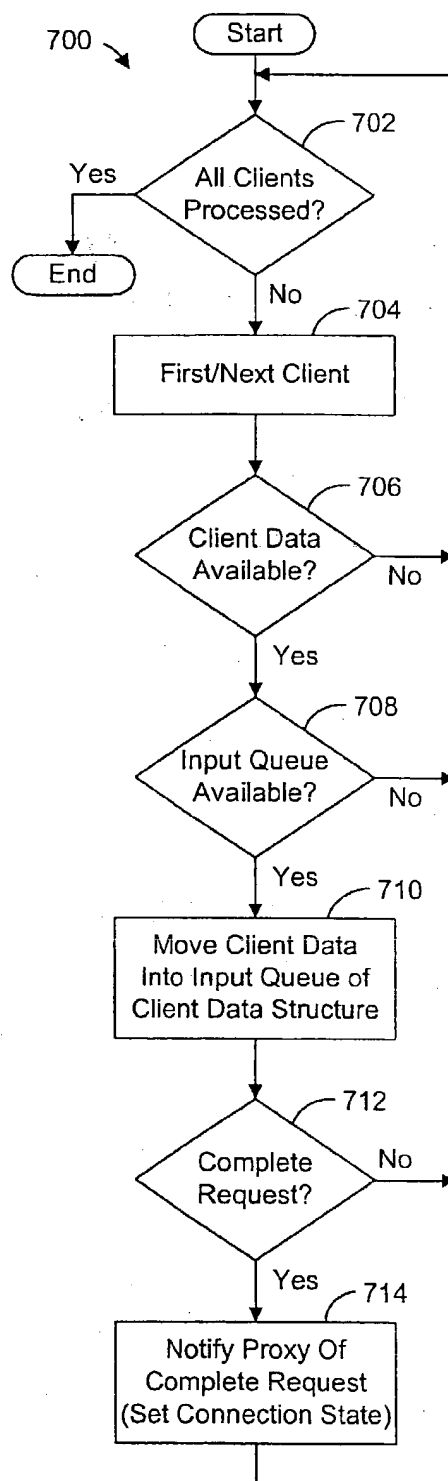


FIG. 7

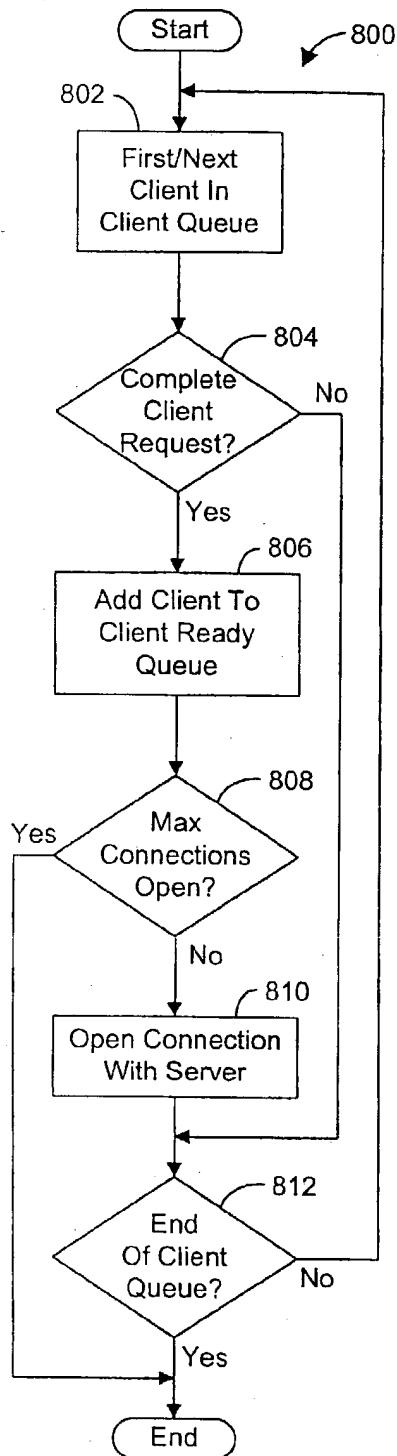


FIG. 8

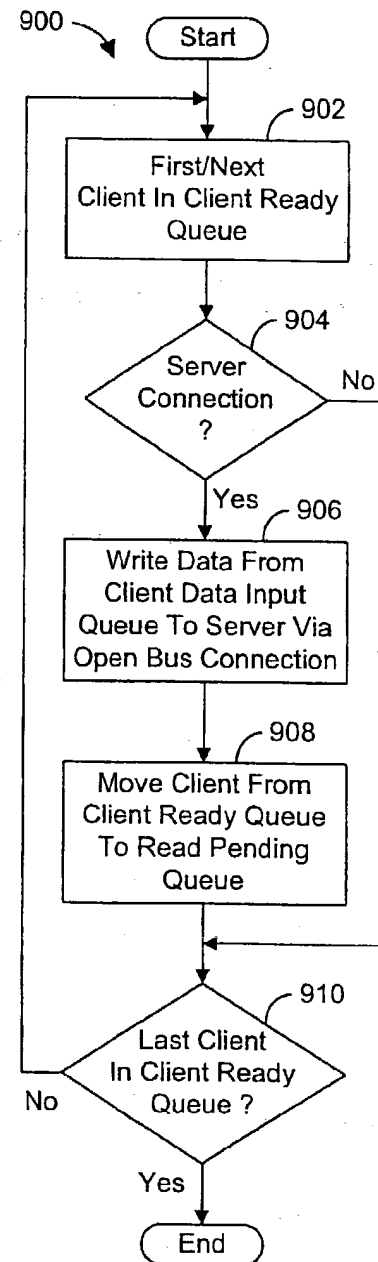
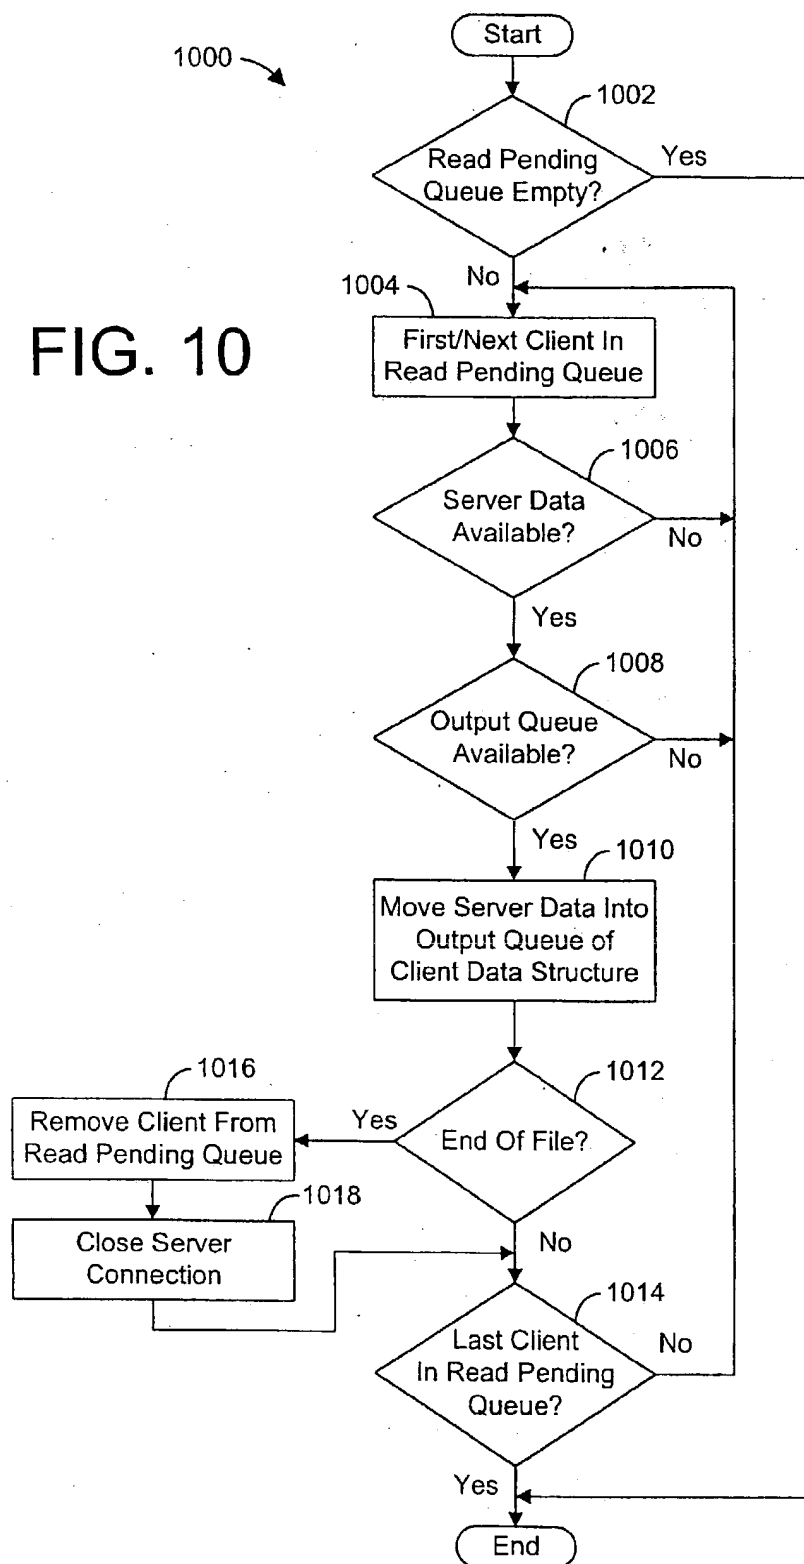


FIG. 9

FIG. 10



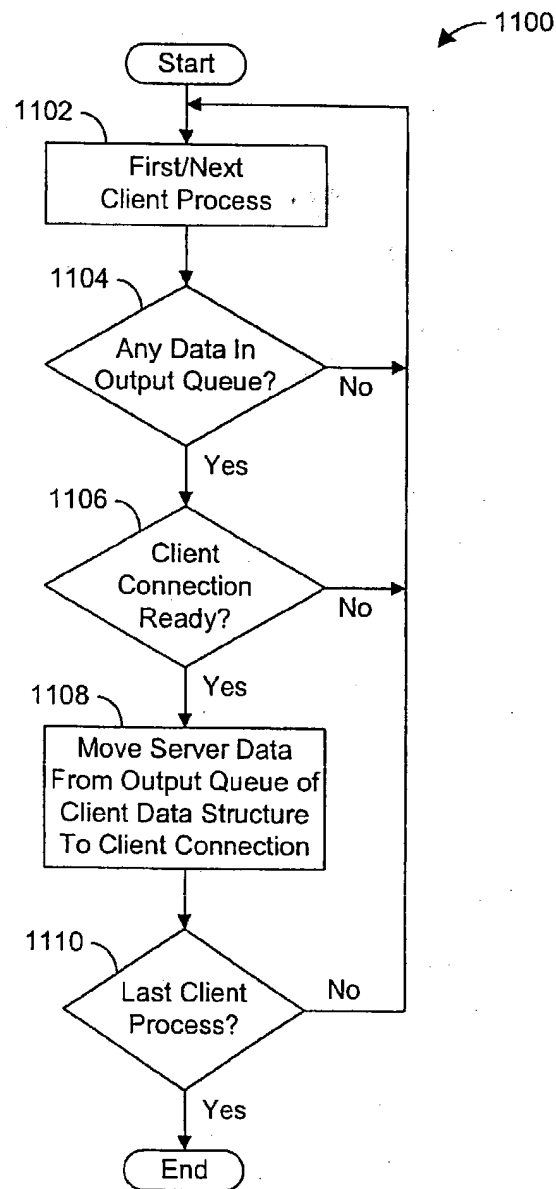


FIG. 11



北京三友知識產權代理有限公司  
Beijing Sanyou Intellectual Property Agency Ltd.

Add: 11F, Block A, Investment Plaza, 27 Jinrong  
Street, Beijing 100032, China  
Tel: 86-10-66211633 86-10-66211624  
Fax: 86-10-66211631 86-10-66211632  
Http://www.san-you.com  
E-mail: info@sanyou.sina.net

## CHARGE MEMORANDUM

RECEIVED  
Original to Ac.

APR 26 2004

GREENBLUM & BERNSTEIN

GREENBLUM & BERNSTEIN, P.L.C.

April 16, 2004

No. 2004651904

Re: Chinese Patent Application No. 021186189

Your Ref: P21857

Our Ref: I02US03520

Report the response to the Official Action concerning the above-identified patent application:

### Services fees

Receiving, docketing and reviewing the first Office Action issued on  
March 19, 2004; studying the specification and the references carefully  
and providing our comments (1 hour X USD 180.00 hourly rate)

USD 180.00

Translating the text of the OA from Chinese into English(2560 characters)

USD 563.20

Subtotal USD 743.20

### Disbursement

Fax and Postage

USD 20.00

Subtotal USD 20.00

Total USD 763.20

Please settle the invoice by TT to:  
Sanyou Intellectual Property Limited  
The Hong Kong and Shanghai Banking Corporation Limited  
1 Queen's Road Central, Hong Kong  
A/C No. 500-115092-838  
SWIFT CODE: HSBC HKHHHKH

Please settle the invoice by cheque to:  
Sanyou Intellectual Property Limited

This Charge Memorandum is due for payment within 30 days of the above date





北京三友知識產權代理有限公司  
Beijing Sanyou Intellectual Property Agency Ltd.

Add: 11F, Block A, Investment Plaza, 27 Jinrong  
Street, Beijing 100032, China  
Tel: 86-10-66211633 86-10-66211624  
Fax: 86-10-66211631 86-10-66211632  
Http://www.san-you.com  
E-mail: info@sanyou.sina.net

## CHARGE MEMORANDUM

**COPY**

GREENBLUM & BERNSTEIN, P.L.C.

April 16, 2004

No. 2004651904

**Re:** Chinese Patent Application No. 021186189

**Your Ref:** P21857

**Our Ref:** I02US03520

Report the response to the Official Action concerning the above-identified patent application:

### Services fees

Receiving, docketing and reviewing the first Office Action issued on  
March 19, 2004; studying the specification and the references carefully  
and providing our comments (1 hour X USD 180.00 hourly rate)

-----USD 180.00

Translating the text of the OA from Chinese into English(2560 characters)

-----USD 563.20

**Subtotal USD 743.20**

### Disbursement

Fax and Postage

-----USD 20.00

**Subtotal USD 20.00**

**Total USD 763.20**

Please settle the invoice by TT to:

Sanyou Intellectual Property Limited

The Hong Kong and Shanghai Banking Corporation Limited

1 Queen's Road Central, Hong Kong

A/C No. 500-115092-838

SWIFT CODE: HSBC HKHHKH

This Charge Memorandum is due for payment within 30 days of the above date.

Please do not abbreviate any words for the name of the company.

Please settle the invoice by cheque to:

Sanyou Intellectual Property Limited